FermiCloud
Testing and Development Platform
Project Plan

Steven Timm

Version 5.1
Last modified 30-Mar-2010

Abstract:

This document describes the program of work for cloud computing testing and development at Fermilab.

Document Revision History:

Version	Date	Author	Comments
1.0	02-Jan-09	Steven Timm	First draft of project plan
2.0	21-Jul-09	Steven Timm	Second draft of project plan
3.0	17-Aug-09	Steven Timm	Version to be presented to CD Management
3.1	18-Aug-09	Steven Timm	After K. Chadwick revisions
3.2	04-Sep-09	Steven Timm	After G. Garzoglio and E. Berman revisions
3.3	18-Sep-09	Steven Timm	After E. Berman revisions
4.0	29-Sep-09	Steven Timm	After first round of stakeholder meetings
5.0	23-Feb-10	Steven Timm	Expand the WBS based on discussions in the meeting
5.1	30-Mar-10	Steven Timm	Further WBS expansion based on previous meeting

Executive Summary:

This project will investigate, design, and deploy a flexible Infrastructure-as-a-service facility ("FermiCloud") for use by the grid and storage developers, integrators, and testers of the Fermilab Computing Division. We are investigating currently-available virtualization and provisioning technologies, discussing requirements with stakeholders, and determining the budget in FTE's and hardware costs. This project will deliver an infrastructure service which deploys just-in-time build and test images for all supported Fermilab operating systems. The virtual images will be operational only as long as they are needed. This will save on floor space and power, and also save significant administrative effort currently used to provision and support legacy hardware. The experience we gain for this project will also be used in the proposed General Physics Computing Facility currently under design for the Intensity Frontier experiments, as well as other stakeholders across the division. We are executing the project in three phases:

- Phase Zero is the planning and requirements gathering phase of the project, detailed in this document.
- Phase One is the deployment of a small cloud on existing hardware as a proof of principle, evaluating various existing virtualization software packages and working with early adopters to find the design that meets the technical and Fermilab security policy requirements.
- Phase Two is the deployment of the production cloud on two racks of new servers as requested in the FermiGrid Tactical Plan, accompanied by a production software stack.

The scope of this project is to build a scientific testing and development private cloud at Fermilab. We will not deploy services on commercial clouds, although we will attempt to find a solution which is compatible with existing commercial clouds. The project does not include virtualization of the classic IT and business functions of the lab, which is under investigation by a different group and quadrant. This project does not include virtualizing the worker nodes of FermiGrid nor submission of virtual machines as jobs to FermiGrid.

Stakeholders:

FermiGrid Services Group (Timm, Chadwick)
Distributed Offline Computing Services Group (Garzoglio, Mhashilkar, Baranovski)
Data Movement + Storage (Oleynik, Perelmutov, Moibenko)
Fermilab Experimental Facilities Dept. (J. Allen, E. Simmonds, L. Ho)
High Performance Parallel Comp. Fac. (Singh)
Comp. Enabling Technologies (Kowalkoski, Paterno)
JDEM (Nielsen)
Virtual Services Group—LSCS quadrant (Rosier)
OSG + LHC Tier 3 Fraser (OSG), Snihur (CMS), Benjamin (ATLAS)
OSG Storage (Levshina)
OSG Gratia project (Canal, Green)
REX Dept (D. Box)
SSE Group (Garren)

Statement of Need:

The initial need for infrastructure-as-a-service comes from the requirements of developers and integrators. The various development and integration groups within the SCF quadrant, including WAN, Data Movement and Storage, Open Science Grid, and FermiGrid, have relied on the FCC1 GridWorks cluster (a.k.a. FAPL cluster) for development and testing. Most of the machines in this cluster are 6 years old or older, and were not in good condition when they were made available to us after being decommissioned from the production farms. Over the past year almost 50% of these machines have broken and been unfixable, causing delays in critical development and testing work. Now we have been forced to turn off all such development machines in FCC1. In addition, there are a number of other development machines around FCC1 and FCC2 which are also aging and have no replacement plans at present. Even if working properly, the current machines have inadequate memory to test grid middleware. The developers and integrators need access to modern hardware, particularly for intense integration and testing cycles. This access is often needed on short notice. An on-demand provisioning system, coupled with the capacity to store the state of the virtual machine when it is not in use, is needed.

Discussions with major stakeholders and groups that manage large amounts of servers have identified three major categories of needs which shape the project. (1) Developers and integrators who need the development, testing and integration resources of FermiCloud itself. (2) The General Physics Computing Facility and other follow-on projects which will depend on the technologies that are tested in this project. (3) Concurrent virtualization-related activities in Fermilab and in the OSG with which we want to cooperate. We summarize them below:

Developers and Integrators:

1. The SRM/dCache developers at Fermilab are heavily dependent on the current GridWorks cluster. They use eight of the existing nodes for a persistent test stand both for internal use and for external SRM developers to test clients against the SRM/dCache server. They need to replace this functionality and also will use the opportunity to have more test machines and integration machines from time to time. The Enstore developers have interest in intermittent use of testing and integration machines as well.
2. Distributed Offline Computing Services Group—The developers in this group are involved with a variety of internal and external projects, all of which need virtual machines for development, testing, and integration. These include but are not limited to Gratia, OSG Storage, SAMGrid, GlideinWMS, MCAS, SAZ, and Software Project Management Infrastructure.
3. Storage Evaluations--G. Garzoglio is currently leading an investigation of current storage use cases on Grid computing at Fermilab and has proposed a program of work which includes further testing of distributed file systems such as Hadoop and Lustre. Such file systems are often closely associated with cloud computing work and an infrastructure-as-a-service platform provides a good testbed to do this work with minimal impact on the sysadmins.
4. LQCD—They are considering distributing virtual machines to remote users for kerberos client and login, also considering the possibility of making 4-processor or 8-processor virtual machines available to users for compilation and testing of small parallel jobs, with the idea that users would work some of the obvious problems out of their job before they get on the main cluster. They have used VirtualBox for their current work.

5. Computing Enabling Technologies—They are doing data workflow management development for various projects including JDEM. These are investigations where we need to simulate a large distributed system with limited resources.

General Physics Computing Facility:

1. General Physics Computing Facility, the proposed home of interactive and small batch analysis for the Intensity Frontier experiments. Virtual machine technology will be used in deploying experiment-specific requirements. It will begin with static virtual machines such as are currently used by FEF or FermiGrid. The program of work of dynamic provisioning in this project will allow a transition to dynamic provisioning in the General Physics Computing Facility. This facility will also include virtual machines with identical configuration to FermiGrid worker nodes so that users can log in and debug their jobs.

Concurrent Virtualization Activities:

1. Server OS Virtualization and provisioning technology. FEF and FGS rely on server virtualization to run large numbers of low-cpu-usage servers on a few machines. FEF has used Virtual Iron with success and is actively investigating its follow-on product, OracleVM. FGS has used open-source methods. Upcoming transitions in the commercial virtualization space and the open-source world make the cloud project a good opportunity to evaluate existing and new solutions. Key to this investigation is identifying products that are reliable and easily manageable.
2. OSG and LHC—There is active effort in the LHC experiments to make a set of distributable virtual machines so that university sites can actively provision a Tier 3 with a minimum of effort. These typically can run under Xen or VMWare. The CERNVM project to create virtual CERN environments on the laptop and desktop is also a significant strategic direction for them.
3. External cloud investigations—Up until now these have included running production (STAR@BNL) and investigation of using cloud for disaster failover (CMS). Several groups are interested in the external cloud as a development platform.
4. Virtual Services Group (LSCS quadrant) is virtualizing a number of servers in the classic IT field using VMWare ESX. They have interest in the cloud due to occasional use cases which come to them that don't warrant using a VMWare license.

Available Technology Investigation:

There are four basic areas of technology investigation identified thus far. (1) Available virtualization hypervisors and the hardware and OS on which they run. In common use at Fermilab thus far are Xen and VMware. RedHat's KVM is available in some test clusters. Microsoft's Hyper-V has been tried for some Windows servers. Commercial as well as open source versions of Xen will be considered. (2) Available virtual machine provisioning and deployment mechanisms. These include open-source emulations of the Amazon Elastic Compute Cloud (EC2) such as Eucalyptus (UCSB) and Nimbus (U of Chi.) They also include Virtual Iron (soon to be incorporated into Oracle VM) and a number of other commercial offerings, including some that RedHat currently has in beta such as oVirt and OpenQRM. (3) Cloud-like object stores capable of storing and retrieving virtual machine input and other files. (4) Highly performant clustered file systems that work under Xen and the hardware to support them which are needed for the VM library and also non-system-image data access. The project will attempt to use existing software systems to the greatest extent possible, but some minor modifications are expected to be necessary to comply with Fermilab's security environment. The initial goals of developer support can be met without a high-performance backend store if necessary.

Requirements:

Requirements collection is still ongoing but we have already collected a large number of requirements.

A. Operating systems that need to be supported:
    1. Client virtual machines: Linux (any versions currently within Fermilab baselines).
    2. Server host machines: Linux (any versions currently within Fermilab baselines).
    3. Desktop host machines: Linux, Windows, Mac.
B. Functionality:
    1. (FGS) Ability to deploy worker node virtual machines to offer batch slots for opportunistic use from FermiGrid when developer demand is low.
    2. (FGS) Emulation of the Amazon EC2 API as well as other evolving open-source API's
C. Estimated number of total virtual images needed by developers (still under rapid revision):
    1. (FGS/DOCS) Estimated demand within Grid department: OSG ReSS, 6 VM's, OSG Storage, 20-30 VM's, Grid Services related maintenance and testing 10 VM's (includes GlideinWMS), Fermilab Grid Storage Evaluation, 30-40 VM's,Gratia testing 12 VM's, SAMGrid maintenance and testing 4-6 VM's, OSG nightly install tests 10 VM's. Science dashboards 12-15 VM's. These are in addition to any virtual machines which are now available on the FermiGrid test nodes, which are virtualized but whose virtual machines are not dynamically provisioned, there are approximately 10 of these at the moment.
    2. (DMS) 15-20 VM's for SRM and dCache deployment.
D. Estimated number of images that need to be deployed simultaneously
    1. About 50% of the above, 40-50 VM's.
E. Typical memory and disk footprint of application
    1. Varies significantly but average is between 1-2 GB of RAM.
F. Potential disk and network performance issues due to virtualization.
    1. (FGS) These are still under investigation but initial studies have shown that paravirtualized Xen, the most common solution, has performance very close to native disk and network access so we do not anticipate issues for the applications mentioned above.
G. Configurability, maintainability, documentation
    1. (FEF) Management GUI is essential, willing to buy commercial product to get it.
    2. (FEF) Amount of installation, configuration, customization on part of admins should be minimal.
    3. (FEF/FGS) Good monitoring functionality for machines that are already up.
    4. (FEF) Reliability must be high, management overhead must be low.
H. Network topology
    1. (FGS) Need plan to serve either dynamic or fixed IP's. Latter are needed for grid services which expect a non-DHCP hostname.
    2. (FGS) Are private VLAN's needed on a customer-by-customer basis? Several other

academic and commercial clouds have used this model successfully. Need to investigate. Some stakeholders have requested private VLAN's.
3. (FGS) We have requested and got a dedicated subnet for these investigations.
4. Ipv6 has been mentioned by several stakeholders as a potential requirement and also a potential way to work with the IP shortage that the explosion of virtual machines will eventually cause, network topology investigation will include investigations of this.

I. Security and patching
   1. (FGS) Need a patching mechanism to fit the OSE baseline, and not let the virtual machine onto the public net until it is patched.

J. VM and OS provisioning requirements
   1. (FGS) Virtual machine provisioning and OS provisioning should be completely automated and supportable by groups which manage large quantities of worker nodes.
   2. (FGS/DOCS) Virtual machines should be able to be paused and saved and restored later on other hardware.
   3. (FGS) Virtual machines should be able to be turned on and off at a fixed time, and also be able to be instantiated to go for a certain length of time.
   4. (FGS/FEF) Important to have live migration and failover of virtual machines in case of hardware failure or scheduled maintenance

K. FS support (DOCS)
   1. Ability to support remotely mounted FS as well as virtual image storage space.

Design and Evaluation:

Based on the requirements collected above, we will conduct the investigations and come up with a design for a full system that meets them. Our current estimates for the FTE effort and equipment costs are described in the FermiGrid Tactical Plan and the Grid Services Tactical Plan. We have already deployed eight worker nodes using static virtualization techniques in GCC Room B. These are currently serving as the OSG Persistent Integration Test Bed. By use of virtualization we have been able to deploy a Condor, PBS and SGE cluster for Integration use in the same space as was currently used only for a Condor cluster before. We have installed 16 more Dell Poweredge 1950 nodes with the SLF5.3 Xen kernel and are beginning tests of dynamic provisioning techniques. These nodes were all initially bought for testing purposes in FY2007. These nodes are all on a new isolated subnet from the production FermiGrid clusters.

The hardware design is complete at this stage. We have ordered 2U dual quad-core servers based on the Intel "Nehalem" Xeon E5540 processor. They will have eight disks apiece, two "system" disks to store the active virtual images and six "data" disks of 2 TB apiece to provide maximum storage space for the storage testing project, served by a high performance RAID card. We will also add Infiniband card to each machine for the purpose of testing MPI under virtualization and to provide a path to future high-performance high-bandwidth storage. The purchase order has been placed for 23 such servers which will give us 184 physical cores and 368 logical cores due to hyperthreading.

Rack space has been identified in GCC room B for the two racks that will be bought in this order. Networking has been supplied with the budget line item to provision the networking for this cluster which will initially at least be one blade on a Cisco 6509 switch.

WAN Group Network Testing:

Due to the high bandwidth requirements of the WAN group testing and the experimental kernel work that they do, some of the newer GridWorks cluster nodes on FCC1 and the existing switch gear will need to remain available for their testing so they can run on bare metal. Their testing is not a good candidate for virtualization and is not intended to be integrated into FermiCloud. It was originally thought that these test nodes would move to GCC room B with FermiCloud but now this group is looking for alternative space in FCC2 and their nodes remain off for now.

WBS:

The following is a summary of the major tasks of the project.

1.	Technology Investigation
1.1	OS and Hypervisor investigation
1.1.1	Base hypervisor technologies evaluation comparing reliability, performance, and interoperability
1.1.1.1	Xen
1.1.1.2	KVM
1.1.1.3	VMWare
1.1.2	Commercial enhanced hypervisor evaluations, comparison of features and functionality
1.1.2.1	Virtual Iron / Oracle VM
1.1.2.2	Citrix XenServer
1.2	Virtual machine deployment and provisioning investigation
1.2.1	Deployment: Cloud software packages investigation
1.2.1.1	Nimbus evaluation
1.2.1.2	Eucalyptus evaluation
1.2.1.3	OpenNebula evaluation
1.2.2	Provisioning: Installing secure VM’s on the cloud
1.2.2.1	Kickstart to create new custom VM’s
1.2.2.2	OS Patching strategy for new, dormant, and restored VM’s
1.2.2.3	Authentication/authorization for VM’s
1.2.3	Scheduling investigation on the cloud
1.3	Object stores and virtual machine library investigation
1.3.1	Virtual machine library design
1.3.2	Object-based cloud storage system evaluation (Amazon S3 emulations)
1.4	Clustered file systems investigation

1.4.1	Lustre performance, bare metal vs. KVM vs. Xen
1.4.2	Hadoop performance, bare metal vs. KVM vs. Xen
1.4.3	Bluearc performance
1.4.4	Xrootd performance, bare metal vs. KVM vs. Xen (time permitting)
1.4.5	Commercial clustered file systems (OCFS2 and others)
1.5	Migration and failover investigation
1.5.1	SAN-based migration and failover
1.5.2	Other migration and failover techniques
1.6	Interoperability/Compatibility specification
1.6.1	Test compatibility of our design with Amazon EC2 and other commercial clouds, using gift research cycles if they are available.
1.6.2	Test compatibility of our design with CERNVM
1.7	Infiniband virtualization evaluation
1.7.1	Correct ConnectX2 drivers in vanilla kernel
1.7.2	PCI pass-thru to single virtual machine, Xen and KVM
1.7.3	Sharing IB card with multiple virtual machines, Xen and KVM
1.7.4	Relative performance metrics of all above methods that work
1.7.5	Survey of storage options available via Infiniband.
1.8	Network topology evaluation
1.8.1	Physical networking design, 1 private + 1 public or bonded public, IP over IB needed?
1.8.2	Investigate leveraging on-site “network jail” and scanning infrastructure for new VM’s
1.8.3	Optimal network configuration—static IP, DHCP, NAT, Ipv6, other?
2.	Design Plan
2.1	Requirements collection from stakeholders and potential early applications
2.2	Hardware technical specification
2.3	Network topology specification
2.4	OS/virtualization technical specification
2.5	VM provisioning/deployment specification
2.6	Object store specification
2.7	Interoperability/Compatibility specification
2.8	Policy specification
3.	Integration and deployment

3.1	Hardware procurement, delivery, installation
3.2	Software acquisition
3.3	Authentication/authorization software modifications if necessary
3.4	User documentation—how to migrate existing VM's and make new ones
4.	Migration of existing test machines to the cloud

Deliverables:

1. Test Cloud Cluster up for investigations:
2. Design plan for production FermiCloud cluster, including finalized hardware and software technical requirements and Statement of Work:
3. Available Technology Investigation report complete:
4. Requirements collection from potential VM users complete:
5. Production Cloud VM provisioning, OS Provisioning, and security patching solution in production:
6. Generalized VM library for users of pre-built virtual machines.
7. Policy draft of cloud computing at FNAL.
8. Transfer of all DOCS Group development and integration work and DMS storage development work to FermiCloud,
9. Integration of FermiCloud technology into other facilities such as General Physics Computing Facility.
10. Documentation and training on how to use the cloud.